

A Parallel Memory-efficient Epistemic Logic Program Solver: Harder, Better, Faster

PATRICK THOR KAHL

Space and Naval Warfare Systems Center Atlantic, North Charleston, SC, USA
patrick.kahl@navy.mil

ANTHONY P. LECLERC

Space and Naval Warfare Systems Center Atlantic, North Charleston, SC, USA
College of Charleston, Charleston, SC, USA
anthony.leclerc@navy.mil / leclerca@cofc.edu

TRAN CAO SON

New Mexico State University, Las Cruces, NM, USA
tson@cs.nmsu.edu

Abstract. As the practical use of answer set programming (ASP) has grown with the development of efficient solvers, we expect a growing interest in extensions of ASP as their semantics stabilize and solvers supporting them mature. Epistemic Specifications, which adds modal operators K and M to the language of ASP, is one such extension. We call a program in this language an epistemic logic program (ELP). Solvers have thus far been practical for only the simplest ELPs due to exponential growth of the search space. We describe a solver that is able to solve harder problems better (e.g., without exponentially-growing memory needs w.r.t. K and M occurrences) and faster than any other known ELP solver.

Keywords: Epistemic Logic Program Solver, Logic Programming Tools, Epistemic Specifications, Answer Set Programming Extensions, Solver Algorithms

1 Introduction

The language of *Epistemic Specifications* [12,13] was introduced in the early 1990s by Michael Gelfond after observing the need for more powerful introspective reasoning than that offered by answer set programming (ASP) alone, extending ASP with modal operators K (“known”) and M (“may be true”). A program written in this language is called an *epistemic logic program* (ELP), with semantics defined using the notion of a *world view*—a collection of sets of literals (*belief sets*), analogous to the answer sets of an ASP program. A renewed interest in Epistemic Specifications [27,8] in 2011 included a proposed change to the semantics by Gelfond [14] in a preliminary effort to avoid unintended world views. That work was continued by Kahl et al. [19,20] in the hopes of finding a satisfactory semantics with respect to intuition and modeling of problems. Later attempts to improve the semantics were offered by Su et al. [26,9] and more recently by Shen & Eiter [24] to further address unintended world view issues.

Along with maturation of the language were various attempts at developing a solver or inference engine, including *ELMO* by Watson [29], *sismodels* by Balduccini [4], *Wviews* by Kelly [28,21] implementing Yan Zhang’s algorithm [30], *ESmodels* by Zhizheng Zhang et al. [25,7,34,33], and most recently *ELPS* by Balai [2,3].

In this paper, we present results of our efforts to implement a new ELP solver that incorporates updated semantics and uses a scalable approach that greatly decreases both the memory and time required for solving harder ELPs compared to other solvers.

2 Syntax and Semantics

In general, the syntax and semantics of the language of Epistemic Specifications follow that of ASP with the notable addition of modal operators K and M, plus the new notion of a *world view*. We assume familiarity with logic programming—ASP in particular. For a good introduction, see any of [6,18,10,15]. For simplicity, the syntax presented does not include certain language features (e.g., aggregates) of the proposed ASP core standard [1], but that does not mean such features should necessarily be considered excluded from the language. As previous authors have done before us, we will use $AS(\mathcal{P})$ to denote the set of all answer sets of an ASP program \mathcal{P} .

2.1 Syntax

An epistemic logic program is a set of rules of the form

$$\ell_1 \text{ or } \dots \text{ or } \ell_k \leftarrow g_1, \dots, g_m, \text{ not } g_{m+1}, \dots, \text{ not } g_n.$$

where $k \geq 0$, $m \geq 0$, $n \geq m$, each ℓ_i is a literal (an atom or a classically-negated atom), and each g_i is either a literal (often called an *objective literal* within the context of Epistemic Specifications), or a *subjective literal* (a literal immediately preceded by K or M). As in ASP, a rule having a literal with a variable term is a shorthand for all ground instantiations of the rule.

2.2 Semantics

It should be noted that the semantics described below differs somewhat from that of [19,20] as a world view has an additional requirement based on [24].

Definition 1. [When a Subjective Literal Is Satisfied]

Let W be a non-empty set of consistent sets of ground literals, and ℓ be a ground literal.

- $K\ell$ is satisfied by W if $\forall A \in W : \ell \in A$.
- $\text{not } K\ell$ is satisfied by W if $\exists A \in W : \ell \notin A$.
- $M\ell$ is satisfied by W if $\exists A \in W : \ell \in A$.
- $\text{not } M\ell$ is satisfied by W if $\forall A \in W : \ell \notin A$.

We will use symbol \models to mean *satisfies*; e.g., $W \models K\ell$ means $K\ell$ is satisfied by W .

Definition 2. [Modal Reduct]

Let Π be a ground epistemic logic program and W be a non-empty set of consistent sets of ground literals. We denote by Π^W the *modal reduct of Π with respect to W* defined as the ASP program¹ obtained from Π by replacing/removing subjective literals and/or deleting associated rules in Π per the following table:

subjective literal φ	if $W \models \varphi$ then...	if $W \not\models \varphi$ then...
$K\ell$	replace $K\ell$ with ℓ	delete rule containing $K\ell$
$\text{not } K\ell$	remove $\text{not } K\ell$	replace $\text{not } K\ell$ with $\text{not } \ell$
$M\ell$	remove $M\ell$	replace $M\ell$ with $\text{not not } \ell$
$\text{not } M\ell$	replace $\text{not } M\ell$ with $\text{not } \ell$	delete rule containing $\text{not } M\ell$

¹ with nested expressions of the form $\text{not not } \ell$ as defined in [22]

Definition 3. [Epistemic Negations²]

Let Π be a ground epistemic logic program and W be a non-empty set of consistent sets of literals. We denote by $E_P(\Pi)$ the set of distinct subjective literals in Π taking the form $\text{not } K\ell$ and $M\ell$ (called *epistemic negations*) as follows:

$$E_P(\Pi) = \{ \text{not } K\ell : K\ell \text{ appears in } \Pi \} \cup \{ M\ell : M\ell \text{ appears in } \Pi \}.$$

We use Φ to denote a subset of $E_P(\Pi)$, and we denote by Φ_W the subset of epistemic negations in $E_P(\Pi)$ that are satisfied by W ; i.e., $\Phi_W = \{ \varphi : \varphi \in E_P(\Pi) \wedge W \models \varphi \}$.

Definition 4. [World View]

Let Π be a ground epistemic logic program and W be a non-empty set of consistent sets of literals. W is a *world view* of Π if:

- $W = \text{AS}(\Pi^W)$; and
- there is no W' such that $W' = \text{AS}(\Pi^{W'})$ and $\Phi_{W'} \supset \Phi_W$.

Note here the addition of a *maximality requirement* on Φ_W with respect to other guesses (corresponding to other candidate world views) that is not in the semantics of [19,20]. See [24] for discussion concerning the intuition behind the proposed new semantics.

2.3 Discussion and Additional Definitions

The semantics of an epistemic logic program Π as described herein is equivalent to the semantics described by Shen & Eiter in [24] for Π translated to their syntax. The proof can be found in the appendix.

Although we prefer our syntax to that proposed in [24], we find the definition of an *epistemic reduct* an excellent tool for describing the new semantics, particularly with the emphasis on our ELP solver. We therefore present the following additional definitions.

Definition 5. [Epistemic Reduct, Reduct-verifiable]

Let Π be a ground epistemic logic program, Φ be a subset of $E_P(\Pi)$, and $\Psi = E_P(\Pi) \setminus \Phi$. We denote by Π^Φ the *epistemic reduct of Π with respect to Φ* defined as the ASP program obtained from Π (assuming existence of corresponding W) by considering as

SATISFIED: the subjective literals in Φ and the complements of the subjective literals in Ψ

NOT SATISFIED: the subjective literals in Ψ and the complements of the subjective literals in Φ ³

and taking actions according to the table given for the *modal reduct* in Definition 2.

For given Π , we say that Φ is *reduct-verifiable* if $W = \text{AS}(\Pi^\Phi)$, $W \neq \emptyset$, and $\Phi_W = \Phi$.

Definition 6. [World View (alternative definition)]

Let Π be a ground ELP, $\Phi \subseteq E_P(\Pi)$, and $W = \text{AS}(\Pi^\Phi)$. W is a *world view* of Π if:

- Φ is reduct-verifiable and
- there exists no reduct-verifiable $\Phi' \subseteq E_P(\Pi)$ such that $\Phi' \supset \Phi$.

² introduced by Shen & Eiter in [24] using a different syntax (see the appendix)

³ $K\ell$ and $\text{not } K\ell$ are complements; $M\ell$ and $\text{not } M\ell$ are complements

3 Algorithms for Computing World Views

Definitions 5 and 6 show that a world view of a program Π can be computed by guessing a set $\Phi \subseteq E_P(\Pi)$ and verifying that Φ is maximal (with respect to \subseteq) reduct-verifiable. As such, we use the term *guess* to refer to a set of elements from $E_P(\Pi)$ in the descriptions of our algorithms. For completeness of the paper, we include herein the basic algorithm implemented in existing state-of-the-art ELP solvers (Algorithm 1). A pictorial description of Algorithm 1 is given in Figure 1(a). It is easy to see Algorithm 1 is complete if **Translation** (Step 2) guarantees all world views of Π are disjoint subsets (groups) of the answer sets of Π' . One such translation was proposed in [20].

Algorithm 1: Compute World Views – Old Algorithm

Input: Π : an epistemic logic program
Output: The set $\Omega = \{\omega : \omega \text{ is a world view of } \Pi\}$

- 1 $\Omega \leftarrow \emptyset$
- 2 **Translation:** ASP program Π' is created from input ELP Π (effecting generation of Π^Φ for all guesses $\Phi \subseteq E_P(\Pi)$).
- 3 **Computation:** Π' is solved using an ASP solver.
- 4 **Aggregation:** Answer sets of Π' are grouped according to corresponding Φ .
- 5 **Verification:** Each group G is verified (if $\Phi_G = \Phi$ then G is added to Ω).
- 6 **return** Ω

Although Algorithm 1 is simple and easy to implement as an add-on to an ASP solver, its main drawback lies in Steps 3 & 4. Specifically, it requires that *all* answer sets of program Π' are computed and then grouped before world views can be identified. This works well for very small programs but can be impractical memory-wise since the number of guesses grows exponentially with the number of subjective literals. In addition, existing implementations of this algorithm do not allow for early termination when the number of world views found reaches a user-specified limit (e.g., 1).

3.1 New Algorithm

Our new algorithm improves Algorithm 1 by addressing the memory growth problem. It divides the work into parts instead of computing everything in one call. The basic steps are given in Algorithm 2. A pictorial description of the algorithm is given in Figure 1(b). In the form shown, the algorithm does not address the maximality requirement of the updated semantics. Details of how the values of guesses are added to the program in the **Partition** step (Step 4) will be described in the implementation section, where selection of guesses is done systematically via an *ordered search* with *pruning* that ensures the maximality requirement is properly addressed. It is easy to see how limiting the number of guesses during the **Partition** step of Algorithm 2 can alleviate the memory concerns of Algorithm 1; however, as the selection of guesses for each iteration is not specified, the algorithm can be implemented in a variety of ways. The algorithm will thus be further refined in Algorithms 3, 4, and 5.

Algorithm 2: Compute World Views – New Algorithm

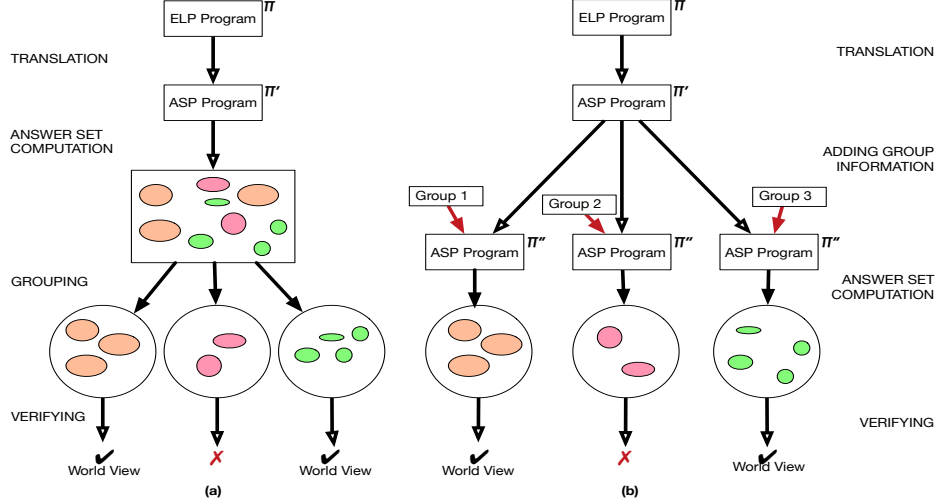
Input: Π : an epistemic logic program
Output: The set $\Omega = \{\omega : \omega \text{ is a world view of } \Pi\}$

- 1 $\Omega \leftarrow \emptyset$
- 2 **Translation:** ASP program Π' is created from input ELP Π (without the inclusion of guesses for subjective literals)
- 3 **repeat**
- 4 **Partition:** Π'' is created by adding ASP representation of guesses (Φ values) to Π'
- 5 **Computation:** Π'' is solved using an ASP solver
- 6 **Aggregation:** Answer sets of Π'' are grouped according to corresponding Φ .
- 7 **Verification:** Each group G is verified (if $\Phi_G = \Phi$ then G is added to Ω).
- 8 **until** all guesses tried
- 9 **return** Ω

Algorithm 2 offers advantages over Algorithm 1 by providing a divide-and-conquer search method that is configurable to allow for:

- efficient use of memory,
- parallel processing, and
- early termination after a pre-specified number of world views found.

Fig. 1. Side-by-side Visual Comparison of the Old (a) and New (b) Algorithms



3.2 Instantiations of the New Algorithm: Overview

The basic idea behind the new algorithm is to follow Definition 6 in computing the world views. To ensure completeness in an efficient manner, guesses need to be selected in a way that facilitates satisfaction of the second condition of the definition. In our approach, guesses are ranked (into *levels*) in decreasing order by their cardinality. The

cardinality of a guess, $|\Phi|$, is called the *guess size*. A guess is filtered out if it is a subset of a prior guess that was found to be reduct-verifiable. Combined with the selection order, this eliminates the possibility of a reduct-verifiable guess being a superset of any guess that is not filtered out. We therefore start with the guess $\Phi = E_P(\Pi)$ and work systematically down in order of guess size, filtering out guesses that are subsets of previously found world views. So to determine if a guess yields a world view, we need only compute the answer sets of the epistemic reduct (Π^Φ) and check that Φ is reduct-verifiable. Algorithm 3 implements this idea. In this algorithm, $\text{ASP}(\Phi)$ denotes the ASP representation of subjective literals that are considered satisfied when creating Π^Φ , and Π' is constructed in such a way that the corresponding $\text{AS}(\Pi'')^4 = \text{AS}(\Pi^\Phi)$.

Algorithm 3: Compute World Views – Level Order, One Guess at a Time

Input: Π : an epistemic logic program, n_Ω : max number of world views desired
Output: The set $\Omega = \{\omega : \omega \text{ is a world view of } \Pi\}$

```

1  $\Omega \leftarrow \emptyset$ 
2  $\Pi' \leftarrow$  result of Step 2 of Algorithm 2
3 foreach subset  $\Phi$  of  $E_P(\Pi)$  in decreasing order of  $|\Phi|$  do
4   if  $(\exists \omega \in \Omega : \Phi \subset \Phi_\omega)$  then continue // pruning filter
5   construct  $\Pi'' = \Pi' \cup \text{ASP}(\Phi)$ 
6   compute  $C = \text{AS}(\Pi'')$  // candidate world view
7   if  $(C \neq \emptyset) \wedge (\Phi_C = \Phi)$  then add  $C$  to  $\Omega$  // check if world view
8   if  $|\Omega| = n_\Omega$  then break // return if desired # world views found
9 return  $\Omega$ 

```

Note the pruning filter on Line 4 of Algorithm 3 that, combined with the search order, implements the maximality requirement of Definition 6. This filter can be turned off to revert to the semantics of [20]. We refer to this filter later in Algorithms 4 & 5 without explicit pseudocode when we say *filtered*.

Although Algorithm 3 solves the memory problem, it increases the number of calls to the ASP solver. As such, the algorithm can be inefficient. To address this concern, we implement Algorithm 4 that computes the world views for multiple guesses at the same time. The trade-off is that grouping (the **Aggregation** step of Algorithm 2) must be performed. Additionally, care needs to be taken to ensure completeness and avoid repetition. We accomplish this by requiring that all guesses in a set have the same guess size, all sets of guesses are pairwise disjoint, and all guesses of the same size are tried (or filtered out) before moving to the next level. This is achieved as follows. We introduce parameter n_G representing the (maximum) number of guesses per ASP solver call, and partition each level $L_k = \{\Phi : (\Phi \subseteq E_P(\Pi)) \wedge (|\Phi| = k)\}$ (for $0 \leq k \leq |E_P(\Pi)|$) into groups of at most n_G guesses. Detail on the implementation is given in the next subsection. Figure 2 shows a visual description of this algorithm.

Finally, in order to utilize the computing power of modern computers and enhance the performance of the solver, we implement a parallel version of the algorithm. For brevity, we omit parameter n_Ω and associated pseudocode for early termination when desired number of world views is found.

⁴ modulo certain fresh literals (see Section 3.3)

Algorithm 4: Compute World Views – Level Order, Multiple Guesses at a Time

Input: Π : an epistemic logic program, n_G : #groups per ASP solver call
Output: The set $\Omega = \{\omega : \omega \text{ is a world view of } \Pi\}$

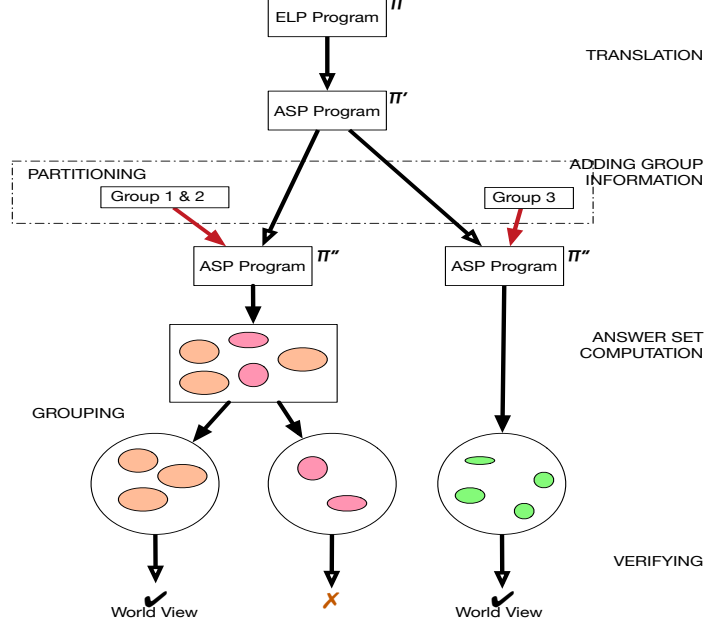
```
1  $\Omega \leftarrow \emptyset$ 
2  $\Pi' \leftarrow$  result of Step 2 of Algorithm 2
3 for  $k \leftarrow |E_P(\Pi)|$  downto 0 do // for each level do ...
4   Partition  $L_k$  into  $t = \left\lceil \frac{|L_k|}{n_G} \right\rceil$  groups  $G_k^1, \dots, G_k^t$  such that  $|G_k^i| \leq n_G$ 
5   foreach filtered group  $G_k^i$  of  $L_k$  do
6     construct  $\Pi'' = \Pi' \cup \mathbf{ASP}(G_k^i)$ 
7     compute  $G^{as} = \mathbf{AS}(\Pi'')$ 
8     foreach associated candidate world view  $C$  of  $G^{as}$  do // group
9       if  $(C \neq \emptyset) \wedge (\Phi_C = \Phi)$  then add  $C$  to  $\Omega$  // check if world view
10 return  $\Omega$ 
```

Algorithm 5: Compute World Views – Parallel Version

Input: Π : an epistemic logic program, n_G : #groups per ASP solver call, n_p : #processes
Output: The set $\Omega = \{\omega : \omega \text{ is a world view of } \Pi\}$

```
1  $\Omega \leftarrow \emptyset$ 
2  $\Pi' \leftarrow$  result of Step 2 of Algorithm 2
3 for  $k \leftarrow |E_P(\Pi)|$  downto 0 do // for each level do ...
4   Partition  $L_k$  into  $t = \left\lceil \frac{|L_k|}{n_G} \right\rceil$  groups  $G_k^1, \dots, G_k^t$  such that  $|G_k^i| \leq n_G$ 
5   Mark every group as not considered
6   repeat
7     for  $i \leftarrow 1$  to  $n_p$  do // for each processor do ...
8       select a filtered group  $G$  from  $G_k^1, \dots, G_k^t$  that is marked as not considered
9       mark group  $G$  as considered
10      construct  $\Pi''_i = \Pi' \cup \mathbf{ASP}(G)$ 
11      for  $i \leftarrow 1$  to  $n_p$  do in parallel // solve in parallel
12        compute  $G_i^{as} = \mathbf{AS}(\Pi''_i)$ 
13        // accumulate the results of each processor and ...
14      for  $i \leftarrow 1$  to  $n_p$  do
15        foreach associated candidate world view  $C$  of  $G_i^{as}$  do // group
16          if  $(C \neq \emptyset) \wedge (\Phi_C = \Phi)$  then add  $C$  to  $\Omega$  // check if world view
17  until every group is either marked as considered or filtered out
18 return  $\Omega$ 
```

Fig. 2. Visual Depiction of Algorithm 4



3.3 Implementation Representation

Let Π be a ground ELP and $E_P(\Pi)$ be its set of epistemic negations. Let \mathcal{I}_{E_P} be an enumeration of $E_P(\Pi)$ and for each $\varphi \in E_P(\Pi)$, let $ord(\varphi)$ be the index of φ in \mathcal{I}_{E_P} . We represent each guess $\Phi \subseteq E_P(\Pi)$ with a bitvector of length $n = |E_P(\Pi)|$ as follows: $B_\Phi = [b_1, b_2, \dots, b_n]$ where for $i = ord(\varphi)$, bit $b_i = 1$ iff $\varphi \in \Phi$.

Each bitvector representing a guess can be represented by an integer X such that $0 \leq X \leq 2^n - 1$; thus, each X represents a unique subset of $E_P(\Pi)$ and vice versa. This one-to-one correspondence is useful in the implementation as it allows for fast checking of the subset relation between guesses (Line 4, Algorithm 3). For later use, we denote by $popcount(X)$ the number of one bits in the binary representation of X .

As an example, consider $E_P(\Pi) = \{\text{not } K \ell_1, \text{not } K \ell_2, M \ell_3\}$. Assuming we enumerate the elements of $E_P(\Pi)$ from left to right, the bitvector $B_\Phi = [b_1, b_2, b_3]$ would be constructed with bit b_1 representing the truth value (w.r.t. membership in Φ , where 1=True and 0=False) of $\text{not } K \ell_1$, b_2 representing the truth value of $\text{not } K \ell_2$, and b_3 representing the truth value of $M \ell_3$. For example, a bitvector $B_\Phi = [0, 1, 1]$ would represent $\Phi = \{\text{not } K \ell_2, M \ell_3\}$.

For the translation of an ELP Π to an ASP program, we use a variant of the one developed in [20]. For each literal ℓ in the epistemic negations of $E_P(\Pi)$ of the form $\text{not } K \ell$, let $k_ \ell$, $k0_ \ell$, and $k1_ \ell$ be fresh atoms created by prefixing ℓ with $k_$, $k0_$, and $k1_$ (respectively), substituting 2 for \neg if ℓ is a classically-negated atom. Likewise, for epistemic negations of the form $M \ell$ in $E_P(\Pi)$, let $m_ \ell$, $m0_ \ell$, $m1_ \ell$ be fresh atoms

created in like fashion. For example, if $\ell = p(a)$ then $k_ \ell$ denotes $k_p(a)$, but if $\ell = \neg p(a)$ then $k_ \ell$ denotes $k_2p(a)$. An atom denoted by $k_ \ell$, $k0_ \ell$, or $k1_ \ell$ will be referred to as a *k-atom*, whereas an atom denoted by $m_ \ell$, $m0_ \ell$, or $m1_ \ell$ will be referred to as an *m-atom*. By a *k-/m-literal* we mean a k-/m-atom or its negation. For a set of sets of literals C , we use the notation $C_{\setminus km}$ to mean C modulo k-/m-literals (i.e., with k-/m-literals removed from the sets of C).

Intuitive meanings of the atoms are: $k1_ \ell$ stands for “K ℓ is True,” $k0_ \ell$ stands for “K ℓ is False,” $m1_ \ell$ stands for “M ℓ is True,” and $m0_ \ell$ stands for “M ℓ is False.” We can thus view $k1_ \ell$ as corresponding to K ℓ , $k0_ \ell$ as corresponding to not K ℓ , etc.

3.4 Implementation Details

We describe here implementation details for Algorithm 5.

1. **Translation:** Given ground ELP Π , we create ASP program Π' by (i) leaving rules without subjective literals unchanged; and (ii) for rules containing subjective literals, replacing subjective literals and adding new rules per the following table.

SUBJECTIVE LITERAL	REPLACE WITH	ADD RULES
K ℓ	not $\neg k_ \ell$, ℓ	$\neg k_ \ell \leftarrow k0_ \ell$.
not K ℓ	$\neg k_ \ell$	$\neg k_ \ell \leftarrow k1_ \ell$, not ℓ .
M ℓ	$m_ \ell$	$m_ \ell \leftarrow m1_ \ell$.
not M ℓ	not $m_ \ell$	$m_ \ell \leftarrow m0_ \ell$, not not ℓ .

We note that this translation is slightly different from the translation in [20] in that it does not add the rules for guessing the values of elements in $E_P(\Pi)$. On the other hand, it implies that if equivalent rules for guessing the values of elements in $E_P(\Pi)$ are added to Π' then correctness of the algorithm is maintained. This property is guaranteed by the partitioning step of the algorithm.

2. **Partition:** The algorithm employs a bitvector representation for guesses in partitioning the search space. The partitioning of a level, L_k , occurs on Line 4 of the algorithm. Each group of bitvectors, G_k^1, \dots, G_k^t for $t = \left\lceil \frac{|L_k|}{n_G} \right\rceil$, is produced “on demand” with group size at most n_G , and G_k^1 containing the first n_G elements of L_k , G_k^2 containing the next n_G elements of L_k , etc. Partitioning is accomplished using a generating function “seeded” with an appropriate bitvector of length n with k one bits. Each subsequent call to the generating function produces the “next” bitvector of length n with k one bits. Thus, storage is required only for representing groups of guesses currently under consideration.
3. $\Pi'' = \Pi' \cup \text{ASP}(G)$, where G is a set of bitvectors, is implemented as $\Pi' \cup \text{ASP}(\{X_B : (B \in G) \wedge (X_B \text{ is the integer whose binary representation is } B)\})$.
4. **Aggregation:** Answer sets of Π'' are grouped by common k-/m-atoms of the form $k0_ \ell$, $k1_ \ell$, $m0_ \ell$, and $m1_ \ell$, each group representing a candidate world view. It is easy to see that a group’s k-/m-atoms correspond to a guess Φ_X for some $X \in G$.

5. **Verification:** For each group C representing a candidate world view, check that the following conditions are met for all its k - m -atoms:

- (a) if $k1_l$ is in the sets of C , then l is in every set of C ;
- (b) if $k0_l$ is in the sets of C , then l is missing from at least one set of C ;
- (c) if $m1_l$ is in the sets of C , then l is in at least one set of C ; and
- (d) if $m0_l$ is in the sets of C , then l is missing from every set of C .

$C_{\setminus km}$ is a *world view* of Π if the conditions above are met.

3.5 Correctness of Algorithm

The correctness of the algorithm follows from the proofs for soundness and completeness given in [20] when one considers that we are simply partitioning the search space into groups of manageable size (rather than generating all possible combinations of subjective literal truth values en masse) and imposing a popcount-level search order. It is easy to see that the maximality requirement of the new semantics is guaranteed by the level-based search order and filtering by considering that:

- if a set of answer sets C corresponding to Φ_C is found that satisfies the verification conditions of Step 5, then the ordered search and the filtering out of any corresponding $\Phi \subset \Phi_C$ ensures that no previous world view W was found corresponding to Φ_W such that $\Phi_W \supset \Phi_C$; and
- if a set of answer sets C corresponding to Φ_C is found to be a world view, then any set C' corresponding to $\Phi_{C'}$ where $\Phi_{C'} \subset \Phi_C$ will be filtered out thereafter.

4 Test Results

We tested an implementation of our algorithm on a Dell™ Precision™ T3500 server with an Intel® Xeon® W3670@3.2GHz with 6 cores and 12 GB RAM running CentOS v6.7 operating system. For the front-end, we used the ELPS solver by Balai [2] to convert an ELPS program (see [3]) to an associated ASP program, grounding it with ASP grounder *gringo* by Kaminski [23] in order to obtain what would be the result of Step 2 of our algorithm for a corresponding ground ELP from the ungrounded ELPS program. We then took the resulting ground ASP program as input to our solver, *ELPsolve*—a loosely-coupled system that uses ASP solver *clingo* by Kaminski & Kaufmann [23].

This method allowed us to directly compare our solver performance with that of ELPS for the same input programs. The following table shows our results, giving the best run time (elapsed time in seconds) observed over the course of testing. A dash (–) in the ELPS column indicates that a runtime error occurred due to insufficient memory.

Π	$ E_P(\Pi) $	ELPS	ELPsolve	Π	$ E_P(\Pi) $	ELPS	ELPsolve
eligible01	2	<1s	<1s	yale1	5	<1s	<1s
eligible06	12	1s	<1s	yale2	8	<1s	<1s
eligible08	16	16s	1s	yale3	11	<1s	<1s
eligible09	18	150s	2s	yale4	14	<1s	<1s
eligible10	20	–	4s	yale5	17	13s	<1s
eligible12	24	–	119s	yale6	26	–	<1s
eligible14	28	–	1667s	yale7	30	–	2s
eligible16	32	–	16124s	yale8	34	–	76s

The `eligibleN` programs are ELPS implementations of the *scholarship eligibility problem* (see [12]) with N corresponding to the number of students. The `yaleN` programs are ELPS implementations of a version of the *Yale shooting problem* (see [17]) that use the *epistemic conformant planning module* defined in [20] with N corresponding to the horizon (number of steps in a plan). For the `yaleN` programs, `ELPsolve` incorporates two heuristics found during analysis of the conformant planning module:

1. reduction of the search space size by a factor of 4 by recognizing that two specific subjective literals must be satisfied in a world view;⁵ and
2. the horizon corresponds to a particular level (guess size) for searching.

Figure 3 shows how run times for solving `eligible12` with `ELPsolve` improve with an increase in the number of processors used. Note that although the 6-core Xeon processor in our test machine is hyperthread-enabled, it is unlikely that there would be any significant speed-up by simply increasing the number of processors assigned to multi-task beyond 6 as our implementation uses 2-thread multi-threading for each task.

Fig. 3. Run Time vs. #Processors (for program `eligible12`)

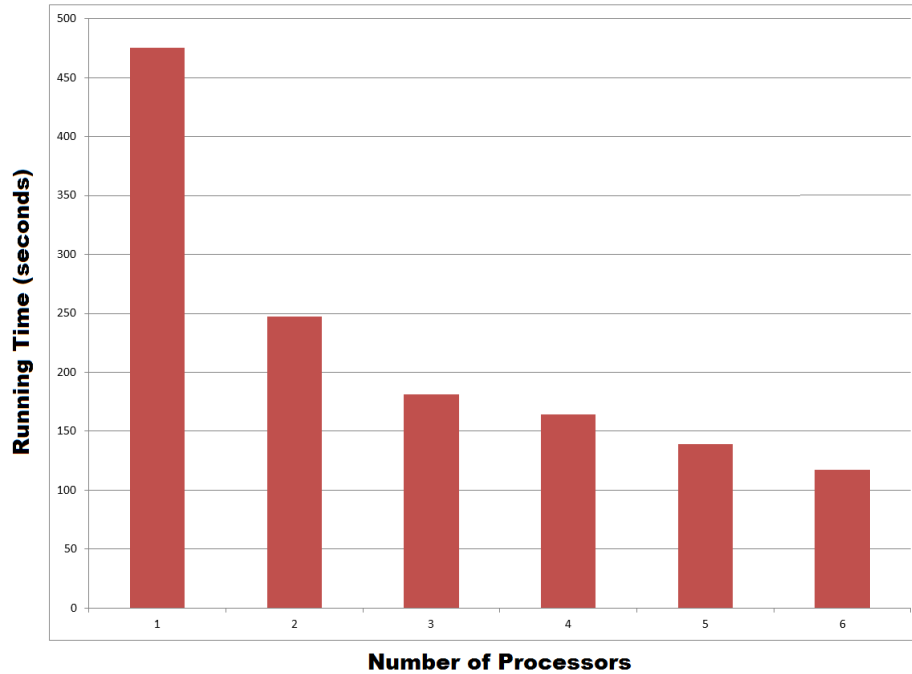
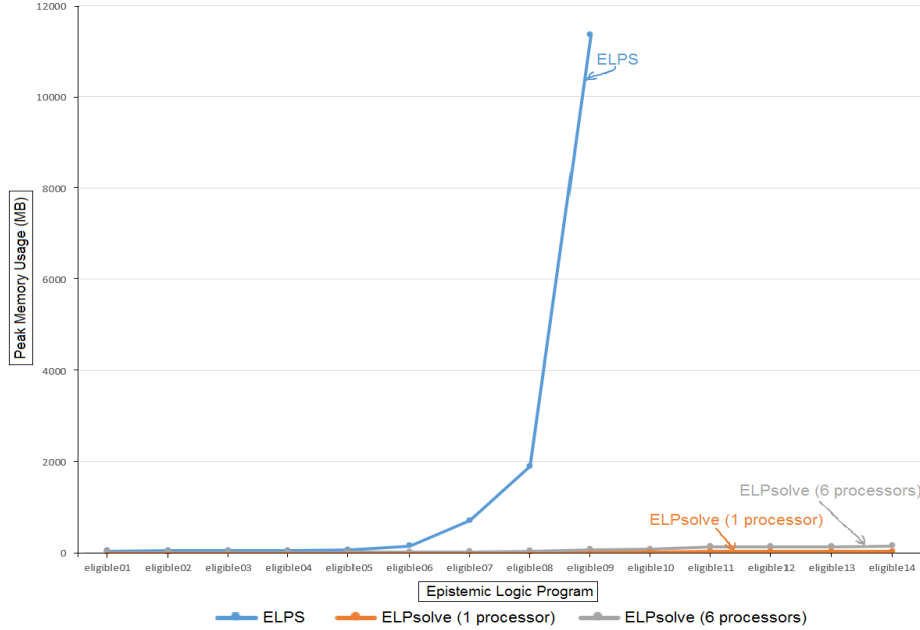


Figure 4 shows the peak memory usage for `ELPS` and `ELPsolve` (max #guesses per ASP call set at 300 for both 1 and 6 processors) when solving various `eligibleN` programs. Note the blue line curving sharply upward for `ELPS` as N grows whereas the memory required for `ELPsolve` remains relatively low and flat.

⁵ This heuristic is generalizable, though a thorough discussion will be left for a future paper.

Fig. 4. Memory Usage (MB): ELPS vs. ELPsolve



5 Related Work

In [30], Yan Zhang investigated computational properties of epistemic logic programs, leading to the development of an algorithm for computing world views. Michael Kelly implemented the algorithm as solver `Wviews` [28], a project for his Honours Thesis.

Cui, Zhizheng Zhang and Zhao [7] investigated the problem of grounding an epistemic logic program. That work culminated in the development of a grounder known as `ESParser`. Their more recent efforts [34,33] include the development of an associated solver called `ESmodels` [25]. See [20] for a comparison of solver performance between `ESmodels` and `ELPS`.

Balai and Kahl [3] extended Epistemic Specifications by adding a sorted signature. A modification of the algorithm in [20,19] was implemented by Balai to develop solver `ELPS` [2]. Our solver complements this work by addressing the memory growth problem and updating it for the new semantics.

Zhizheng Zhang and Shutao Zhang [32] investigated combining ideas from Graded Modal Logic with ASP. Their language can be viewed as an extension of Epistemic Specifications, adding the modal concepts “at least as many as” and “at most as many as” to the language. They continued this work with Wang [31], referring to the language as `GI-log`. A `GI-log` solver called `GISolver` [25] was developed using a generate-and-test algorithm similar (at a high level) to the one described in Algorithm 1.

Shen and Eiter [24] proposed the updated semantics used in this paper and implemented in our solver, albeit using different syntactic notation. Their work resolved most of the open questions raised in [19] and provided inspiration for our algorithm.

6 Conclusions and Future Work

We improved the algorithm and developed a solver for epistemic logic programs that:

- incorporates the latest semantics (as proposed by Shen & Eiter [24]),
- addresses the memory issues that plague some other implementations,
- uses multi-processing to improve performance,
- allows early termination when desired number of solutions found,
- solves harder programs faster, and
- permits solving programs on a typical laptop computer that were previously beyond the capabilities of other solvers with reasonable resources.

For the future, we plan to enhance our solver further by allowing the use of distributed processing to make it reasonable to solve more programs. We also plan to look at generalizing certain heuristics that could significantly reduce the search space for certain classes of programs. Finally, we plan to look at techniques such as *multi-shot solving* [11] to avoid the need for repeatedly restarting the ASP solver from scratch.

We observe that with small modification our implementation could solve GI-log programs [31], though a suitable front end is needed to handle the program syntax.

Follow-up work will be to explore applications for the language as solver improvements make its use more practical. Promising areas include planning & scheduling, policy management, diagnostics, and computer-assisted decision making.

Acknowledgments

We would like to acknowledge the efforts of Yi-Dong Shen and Thomas Eiter whose work provided a positive influence on the development of our solver. We thank Evgenii Balai for providing the front end to our solver by adding an extra option to `ELPS` to output the ASP translation. We couldn't produce our results without the use of `gringo` & `clingo` from the Potassco team, especially Roland Kaminski and Ben Kaufmann. Finally, we thank the anonymous reviewers for their comments and suggestions.

References

1. ASP Standardization Working Group: ASP-Core-2 input language format, version 2.01c. (2013), available at <https://www.mat.unical.it/aspcomp2013/ASPStandardization>
2. Balai, E.: `ELPS` (2015), Texas Tech University, software & documentation available for download at <https://github.com/iensen/elps/wiki/>
3. Balai, E., Kahl, P.: Epistemic logic programs with sorts. In: Inclezan, D., Maratea, M. (eds.) *ASPOCP 2014* (2014)
4. Balduccini, M.: `sismodels` (2001), see <http://www.mbal.tk/> for more information
5. Balduccini, M., Son, T.C. (eds.): *Logic Programming, Knowledge Representation, and Non-monotonic Reasoning—Essays Dedicated to Michael Gelfond on the Occasion of His 65th Birthday*, Lecture Notes in Computer Science, vol. 6565. Springer (2011)
6. Baral, C.: *Knowledge Representation, Reasoning, and Declarative Problem Solving*. Cambridge University Press (2003)

7. Cui, R., Zhang, Z., Zhao, K.: ESParser: An epistemic specification grounder. In: Delgrande, J.P., Faber, W. (eds.) CSSS-12. pp. 1823–1827. IEEE Computer Society CPS (2012)
8. Faber, W., Woltran, S.: Manifold answer-set programs and their applications. In: Balduccini and Son [5], pp. 44–63
9. Fariñas del Cerro, L., Herzig, A., Su, E.I.: Epistemic equilibrium logic. In: Yang, Q., Wooldridge, M. (eds.) IJCAI 2015. AAAI Press / IJCAI (2015)
10. Gebser, M., Kaminski, R., Kaufmann, B., Schaub, T.: Answer Set Solving in Practice. Synthesis Lectures on Artificial Intelligence and Machine Learning, Morgan & Claypool Publishers (2012)
11. Gebser, M., Kaminski, R., Kaufmann, B., Schaub, T.: Clingo = ASP + control: Preliminary report. In: Kowalski, R.A., Bowen, K.A. (eds.) ICLP 2014. vol. abs/1405.3694. The MIT Press (2014)
12. Gelfond, M.: Strong introspection. In: Dean, T.L., McKeown, K. (eds.) AAAI-91. vol. 1, pp. 386–391. AAAI Press / The MIT Press (1991)
13. Gelfond, M.: Logic programming and reasoning with incomplete information. *Annals of Mathematics and Artificial Intelligence* 12(1–2), 89–116 (1994)
14. Gelfond, M.: New definition of epistemic specifications. In: Delgrande, J.P., Faber, W. (eds.) LPNMR-11. Lecture Notes in Computer Science, vol. 6645, pp. 260–265. Springer (2011)
15. Gelfond, M., Kahl, Y.: Knowledge Representation, Reasoning, and the Design of Intelligent Agents: The Answer-Set Programming Approach. Cambridge University Press (2014)
16. Gelfond, M., Lifschitz, V.: Classical negation in logic programs and disjunctive databases. *New Generation Computing* 9(3/4), 365–385 (1991)
17. Hanks, S., McDermott, D.: Nonmonotonic logic and temporal projection. *Artificial Intelligence* 33(3), 379–412 (November 1987)
18. van Harmelen, F., Lifschitz, V., Porter, B. (eds.): Handbook of Knowledge Representation. Foundations of Artificial Intelligence, Elsevier (2008)
19. Kahl, P.: Refining the Semantics for Epistemic Logic Programs. Ph.D. thesis, Texas Tech University, Lubbock, TX, USA (May 2014)
20. Kahl, P., Watson, R., Balai, E., Gelfond, M., Zhang, Y.: The language of epistemic specifications (refined) including a prototype solver. *Journal of Logic and Computation* (2015)
21. Kelly, M.: Wviews: A Worldview Solver for Epistemic Logic Programs. Honour’s thesis, University of Western Sydney (2007)
22. Lifschitz, V., Tang, L.R., Turner, H.: Nested expressions in logic programs. *Annals of Mathematics and Artificial Intelligence* 25, 369–389 (1999)
23. Potassco: `clingo`, `gringo`, `clasp` (2016), University of Potsdam, software & documentation available for download at <http://potassco.sourceforge.net/>
24. Shen, Y.D., Eiter, T.: Evaluating epistemic negation in answer set programming. *Artificial Intelligence* 237, 115–135 (2016)
25. Southeast University: `ESparger`, `ESmodels`, `GISolver` (2015), software & documentation available for download at http://cse.seu.edu.cn/people/seu_zzz/index.htm
26. Su, E.I.: Extensions of Equilibrium Logic by Modal Concepts. Ph.D. thesis, University of Toulouse, Toulouse, France (March 2015)
27. Truszczyński, M.: Revisiting epistemic specifications. In: Balduccini and Son [5], pp. 315–333
28. University of Western Sydney AI Research Group: `Wviews` (2007), UWS AI Research Group (Intelligent Systems Laboratory), software & documentation available for download at <http://staff.scm.uws.edu.au/%7eyan/Wviews.html>
29. Watson, R.G.: An Inference Engine for Epistemic Specifications. Master’s thesis, University of Texas at El Paso (1994)
30. Zhang, Y.: Computational properties of epistemic logic programs. In: Doherty, P., Mylopoulos, J., Welty, C.A. (eds.) KR-06. pp. 308–317. AAAI Press (2006)

31. Zhang, Z., Wang, B., Zhang, S.: Logic programming with graded introspection. In: Incezan, D., Maratea, M. (eds.) ASPOCP 2015 (2015)
32. Zhang, Z., Zhang, S.: Logic programming with graded modality. In: Calimeri, F., Ianni, G., Truszczyński, M. (eds.) LPNMR 2015. Lecture Notes in Artificial Intelligence, vol. 9345. Springer-Verlag (2015)
33. Zhang, Z., Zhao, K.: ESmodels: An epistemic specification solver. CoRR abs/1405.3486 (2014)
34. Zhang, Z., Zhao, K., Cui, R.: ESmodels: An inference engine of epistemic specifications. In: Luo, J. (ed.) ICTAI 2013. pp. 769–774. IEEE (2013)

Appendix: Proof of Semantic Equivalence to Shen-Eiter Semantics

In this appendix, we prove the equivalence between the semantics of epistemic logic programs (Section 2) and the semantics of logic programs with epistemic negation as defined by Shen & Eiter in [24].

For the benefit of the reader, the following table summarizes the notational differences between our language and that of Shen & Eiter for semantically equivalent forms.

OUR NOTATION	SE NOTATION	TERMINOLOGY
\neg	\sim	classical (strong) negation
not	\neg	default negation
not K	not	epistemic negation
K	\neg not	
M	not \neg	
not M	\neg not \neg	

Let Π be a ground epistemic logic program. Let r be a rule in Π . In this appendix, we use the following short hand to represent r :

$$H \leftarrow B^+, \text{not } B^-, \text{K } L_1, \text{not K } L_2, \text{M } L_3, \text{not M } L_4 \quad (1)$$

where H is h_1 or \dots or h_n and $B^+, B^-, L_1, \dots, L_4$ are sets of objective literals. For $L = \{\ell_1, \dots, \ell_m\}$, xL stands for $\{x\ell_1, \dots, x\ell_m\}$ with $x \in \{\text{not}, \text{K}, \text{not K}, \text{M}, \text{not M}\}$. In Shen-Eiter notation [24], the above rule will become the following:

$$H \leftarrow B^+, \neg B^-, \neg \text{not } L_1, \text{not } L_2, \text{not } \neg L_3, \neg \text{not } \neg L_4 \quad (2)$$

We denote by $SE(x)$ the corresponding syntactic form of x in Shen-Eiter notation; e.g., $SE(\Pi)$ denotes the program obtained by translating Π to Shen-Eiter notation. We further define $E_P(SE(\Pi))$ as being equivalent to $SE(E_P(\Pi))$ (see Definition 3). When clear from the context, we may use Φ_W to mean $SE(\Phi_W)$, r to mean $SE(r)$, etc.

In proving the correspondence between the two semantics, we make the following assumptions and notes:

- The presence of classical negation can be eliminated using standard transformation (e.g., the one proposed in [16]). For this reason, we will assume that programs in consideration do not contain classical negation.
- Shen-Eiter programs allow **not** $\neg \ell$ (not K not ℓ in our notation) which we consider equivalent to M ℓ in our syntax.
- In Shen-Eiter programs, $\neg \neg \ell$ (not not ℓ in our notation) is treated as ℓ .

We use “ X -reduct” where X is “KLS” (to refer to the modal reduct discussed in this paper) or “SE” (to refer to the epistemic reduct proposed by Shen & Eiter in [24]). For ASP programs with nested default negation, a “nested-negation removal transformation” refers to the usual substitution of not not ℓ with not ℓ' plus the addition of rule $\ell' \leftarrow \text{not } \ell$ where ℓ' is a fresh literal and answer sets are modulo ℓ' as described in [22]. We use “Gelfond-Lifschitz transformation” to refer to the reduction from logic programs with default negation to positive programs as per [16]. We say that a rule whose body contains a conjunct of the form $\neg \top$ in Shen-Eiter notation is a *useless* rule as this rule is always satisfied (body is false) and cannot be used to justify anything.

In the following, whenever we refer to a rule r , we mean a rule with sets of atoms L_1, \dots, L_4 as in (1). We use Π to denote an arbitrary but fixed ground epistemic logic program. Furthermore, for a program Π , a set of sets of atoms W , a set $A \in W$, and a rule $r \in \Pi$, by the *reduct of r in $(\Pi^W)^A$* (resp. *reduct of $SE(r)$ in $(SE(\Pi)^{\Phi_W})^A$*) we mean the rule (if it exists) obtained from r after:

1. computing the modal reduct of Π with respect to W according to Definition 2 (resp. the epistemic reduct of $SE(\Pi)$ with respect to Φ_W as per [24]); and then
2. computing the Gelfond-Lifschitz transformation on the result with respect to A .

Lemma 1. *Let r be a rule in Π , W be a collection of sets of atoms in Π , and $A \in W$. If there exists some $\ell \in L_1$ of r such that $W \not\models K\ell$ then there is no reduct of r in $(\Pi^W)^A$ and either there is no reduct of r in $(SE(\Pi)^{\Phi_W})^A$ or the reduct of r is a useless rule.*

Proof. $W \not\models K\ell$ implies that r is removed in the KLS-reduct and hence $(\Pi^W)^A$ does not contain the reduct of r .

$W \not\models K\ell$ means that $W \models \text{not } K\ell$. This means that $\text{not } \ell \in \Phi_W$. Thus, the SE-reduct of r in $SE(\Pi)$ is a rule whose body contains $\neg\top$, a useless rule. This implies that $(SE(\Pi)^{\Phi_W})^A$ contains no reduct of r if the Gelfond-Lifschitz transformation removes it; or a useless rule, which is the reduct of r . \square

Lemma 2. *Let r be a rule in Π , W be a collection of sets of atoms in Π , and $A \in W$. If there exists some $\ell \in L_2 \cap A$ such that $W \not\models \text{not } K\ell$, then $(\Pi^W)^A$ and $(SE(\Pi)^{\Phi_W})^A$ contain no reduct of r .*

Proof. $W \not\models \text{not } K\ell$ implies that there exists some $S \in W$ such that $\ell \notin S$. This implies that $W \not\models \text{not } \ell$ and so $\text{not } \ell \notin \Phi_W$.

Let r' and r'' be the result obtained by applying the KLS-reduct and the SE-reduct on r , respectively. We have $r' \in \Pi^W$ and $r'' \in SE(\Pi)^{\Phi_W}$.

Since $W \not\models \text{not } K\ell$, $\text{not } \ell$ occurs in the body of r' . Because $\ell \in A$, the Gelfond-Lifschitz transformation will remove the rule r' from Π^W when constructing $(\Pi^W)^A$, i.e., $(\Pi^W)^A$ does not contain the reduct of r .

Similarly, since $\text{not } \ell \notin \Phi_W$, $\neg\ell$ occurs in the body of r'' in $SE(\Pi)$, and $\ell \in A$, the Gelfond-Lifschitz transformation will remove the rule r'' when constructing $(SE(\Pi)^{\Phi_W})^A$, i.e., $(SE(\Pi)^{\Phi_W})^A$ does not contain the reduct of r .

So, in this case, neither $(\Pi^W)^A$ nor $(SE(\Pi)^{\Phi_W})^A$ contains the reduct of r . \square

Lemma 3. *Let r be a rule in Π , W be a collection of sets of atoms in Π and $A \in W$. If there exists some $\ell \in L_3$ such that $W \not\models M\ell$ then $(\Pi^W)^A$ and $(SE(\Pi)^{\Phi_W})^A$ contain no reduct of r or a useless reduct of r with respect to A .*

Proof. $W \not\models M\ell$ implies that there exists no $S \in W$ such that $\ell \in S$, i.e., for every $S \in W$, $\ell \notin S$. This implies that $W \not\models \text{not } \neg\ell$ and so $\text{not } \neg\ell \notin \Phi_W$. This also implies that $\ell \notin A$.

In this case, we observe that $\text{not not } \ell$ occurs in the body of the KLS-reduct r' of r in Π^W and $\neg \neg \ell$ occurs in the body of the SE-reduct r'' of r in $SE(\Pi)^{\Phi_W}$. Shen & Eiter treat $\neg \neg \ell$ as ℓ . Since $\ell \notin A$, $(SE(\Pi)^{\Phi_W})^A$ will contain no reduct of r (if it is removed by the Gelfond-Lifschitz reduct) or a useless rule with respect to A .

In KLS, nested default negation is treated differently. The nested-negation removal transformation results in $\text{not not } \ell$ appearing in the body of r' . Furthermore, Π^W contains the rule “ $\text{not } \ell \leftarrow \text{not } \ell$ ” where $\text{not } \ell$ is a fresh atom representing $\text{not } \ell$.

Let Π^n be the programs obtained from Π^W after the nested-negation removal transformation. Because $\ell \notin A$, any answer set A' of Π^n whose reduction results into A must contain $\text{not } \ell$. This implies that the Gelfond-Lifschitz transformation will remove r' from Π^n when constructing $(\Pi^W)^A$.

So, in this case, $(\Pi^W)^A$ does not contain the reduct of r and $(SE(\Pi)^{\Phi_W})^A$ contain no reduct of r or a useless reduct of r with respect to A . \square

Lemma 4. *Let r be a rule in Π , W be a collection of sets of atoms in Π and $A \in W$. If there exists some $\ell \in L_4$ such that $W \not\models \text{not } M \ell$ then $(\Pi^W)^A$ contains no reduct of r and $(SE(\Pi)^{\Phi_W})^A$ contains no reduct of r or a useless reduct of r .*

Proof. $W \not\models \text{not } M \ell$ implies that r is removed in the KLS-reduct and hence $(\Pi^W)^A$ does not contain the reduct of r .

$W \not\models \text{not } M \ell$ implies that there exists some $S \in W$ such that $\ell \in S$. It means that $W \not\models K \text{ not } \ell$ and hence $\text{not } \neg \ell \in \Phi_W$. Thus, the SE-reduct of r in $SE(\Pi)^{\Phi_W}$ is a rule whose body contains $\neg \top$. Again, this means that $(SE(\Pi)^{\Phi_W})^A$ contains no reduct of r or a useless rule which is the reduct of r . \square

Theorem 1. *Let W be a collection of sets of atoms in Π . W is a world view of Π under the KLS semantics if and only if W is a candidate world view of $SE(\Pi)$ w.r.t. Φ_W .*

Proof. Observe that for each $\text{not } \ell$ occurring in $SE(\Pi)$, it holds that

- if $\text{not } \ell \in \Phi_W$ then $\text{not } \ell$ is true in W ; or
- if $\text{not } \ell \in E_P(SE(\Pi)) \setminus \Phi_W$ then $\text{not } \ell$ is false in W .

Let $L_2^t = \{\ell \mid \ell \in L_2 \wedge W \models \text{not } K \ell\}$ and $L_2^f = \{\ell \mid \ell \in L_2 \wedge W \not\models \text{not } K \ell\}$. Let $L_3^t = \{\ell \mid \ell \in L_3 \wedge W \models M \ell\}$ and $L_3^f = \{\ell \mid \ell \in L_3 \wedge W \not\models M \ell\}$. Consider $A \in W$ and let r be a rule of the form (1). Lemmas 1-4 show that if $(\Pi^W)^A$ or $(SE(\Pi)^{\Phi_W})^A$ contains a useful reduct of r then the following conditions hold:

- $L_1 \subseteq A$ (Lemma 1).
- for every $\ell \in L_2^f$, $\ell \notin A$ (Lemma 2).
- $L_3^f = \emptyset$ (Lemma 3).
- $L_4 \cap A = \emptyset$ (Lemma 4).
- $B^- \cap A = \emptyset$ (otherwise, the rule will be removed by the Gelfond-Lifschitz transformation).

Under the above conditions, the reduct of r in Π^W is

$$r' : H \leftarrow B^+, \text{not } B^-, L_1, \text{not } L_2^f, \text{not } L_4$$

and its reduct in $(SE(\Pi))^{\Phi_W}$ is

$$r'' : H \leftarrow B^+, \text{not } B^-, \neg\neg L_1, \neg L_2^f, \neg\neg\neg L_4$$

Since $\neg\neg L = L$ as stated in Shen & Eiter's proposal and because $L_4 \cap A = \emptyset$ and $L_2^f \cap A = \emptyset$, we can conclude that the reduct of r in $(\Pi^W)^A$ and in $(SE(\Pi))^{\Phi_W}_A$ is

$$r^+ : \ell_1 \text{ or } \ell_2 \dots \text{ or } \ell_n \leftarrow B^+, L_1$$

In other words, we have that $(\Pi^W)^A$ is identical⁶ to $(SE(\Pi))^{\Phi_W}_A$. As such, if W is a world view of Π , A is an answer set of $(\Pi^W)^A$, i.e., A is also an answer set of $(SE(\Pi))^{\Phi_W}_A$. Conversely, if W is a candidate world view of Π with respect to Φ_W then A is an answer set of $(SE(\Pi))^{\Phi_W}_A$ and hence also an answer set of $(\Pi^W)^A$. Since this holds for every $A \in W$, we have the conclusion of the theorem. \square

⁶ with the exception of the useless rules in different programs and some extra fresh atoms such as *not_a* representing $\text{not } a$ and the rule $\text{not_a} \leftarrow \text{not } a$ introduced to deal with $\text{not not } a$; such atoms can be eliminated by using the splitting set theorem and the fact that, by construction, $\text{not not } a$ appears in Π^W only if $a \notin A$